

What is claimed is:

1 1. A machine-implemented method comprising the steps of:  
2 generating, based on a first set of data associated with a plurality of dimensions, a second  
3 set of data that is denser than the first set of data relative to a first dimension of  
4 the plurality of dimensions;  
5 wherein the first set of data includes a plurality of subsets of data; and  
6 wherein the step of generating includes performing an outer join between each of the  
7 subsets of data and a third set of data.

1 2. The method of claim 1, wherein the first set of data includes rows that are  
2 associated with dimension value combinations in which the dimension value  
3 combinations are combinations of dimension values selected from the plurality of  
4 dimensions,  
5 wherein the second set of data includes corresponding rows for the dimension value  
6 combinations that correspond to the rows of the first set of data,  
7 wherein the corresponding rows are associated with the dimension value combinations,  
8 and  
9 wherein the step of generating comprises the steps of  
10 checking if a corresponding row exists in the second set of data for a set of  
11 dimension value combinations, wherein the set of dimension value  
12 combinations is dense with respect to one dimension; and  
13 creating the row if the corresponding row does not exist.

- 1    3.    The method of claim 2, wherein the step of checking is performed within a set of  
2        nested loop instructions that perform one loop for each dimension value  
3        combination of the set of dimension value combinations.
  
- 1    4.    The method of claim 1, wherein each of the subsets of data is a single row of data.
  
- 1    5.    The method of claim 1, wherein each of the subsets of data is a partition of the  
2        first set data, and is associated with a single dimension value selected from one  
3        dimension of the plurality of dimensions.
  
- 1    6.    The method of claim 1, wherein the step of generating is performed in response to  
2        detecting a data manipulation language statement.
  
- 1    7.    The method of claim 1, wherein the step of generating includes performing the  
2        outer join on a first subset using a first processor and performing the outer join on a  
3        second subset using a second processor that is different than the first processor.
  
- 1    8.    The method of claim 7, wherein the outer join is a right outer join.
  
- 1    9.    The method of claim 8, wherein the outer join is a left outer join.
  
- 1    10.   The method of claim 1, wherein the step of generating is performed by an SQL  
2        engine.

1 11. The method of claim 1, wherein the step of generating includes receiving an  
2 expression that indicates a partitioning key for partitioning the first set of data.

1 12. The method of claim 1, wherein the outer join is associated with join conditions  
2 that includes a Boolean expression.

1 13. The method of claim 1, wherein said first set of data is includes a first set of rows;  
2 and  
3 wherein said outer join is between said first set of rows and a second set of rows, and the  
4 step of generating includes sending each of a plurality of processes a subset of  
5 said first set of rows and all of said second set of rows.

1 14. The method of claim 13 wherein the generating includes  
2 specifying at least one dimension of the plurality of dimensions, and  
3 hash partitioning the first set of data with respect to the dimension specified.

1 15. The method of claim 1, further comprising:  
2 detecting a construct that includes a condition limiting which dimension value  
3 combinations are included in the second set of data; and  
4 in response to detecting the other construct, performing the operation only with respect to  
5 the dimension value combinations to which the second set of data was limited.

1 16. The method of claim 1, wherein the first set of data is associated with a plurality  
2 of dimensions, the second set is associated with the plurality of dimensions, and the  
3 second set of data is denser with respect to one of the plurality of dimensions.

1 17. A machine-implemented method comprising:  
2 generating, based on a first set of data associated with a plurality of dimensions, a second  
3 set of data that is denser than the first set of data relative to a first dimension of  
4 the plurality of dimensions;  
5 wherein the generating is performed without performing a combination of  
6 a sort of the first set of data for distinct values of a second dimension of the  
7 plurality of dimensions,  
8 generating a first set of rows by performing a cross product of the distinct values  
9 found and a set of dimension values of the first dimension, and  
10 adding to said first set of rows, rows corresponding to dimension values of the set  
11 of the dimension values for which no row exists in the first set of rows.

1 18. The machine-implemented method of claim 17, wherein the generating is performed  
2 without performing a sort of the first set of data in which the sort of the first set of  
3 data is used to find distinct values of a second dimension of the plurality of  
4 dimensions.

1 19. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 1.

1 20. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 2.

1 21. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 3.

1 22. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 4.

1 23. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 5.

1 24. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 6.

1 25. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 7.

1 26. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 8.

1 27. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 9.

1 28. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 10.

1 29. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 11.

1 30. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 12.

1 31. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 13.

1 32. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 14.

1 33. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 15.

1 34. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 16.

1 35. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 17.

1 36. A machine-readable medium carrying one or more sequences of instructions, which  
2 when executed by one or more processors, causes the one or more processors to perform the  
3 method recited in Claim 18.

1 37. A system comprising:  
2 one or more processors; and  
3 a machine-readable medium carrying one or more sequences of instructions, which when  
4 executed by the one or more processors, causes the one or more processors to perform  
5 the method recited in Claim 1.

1 38. A system comprising:  
2 one or more processors; and

3 a machine-readable medium carrying one or more sequences of instructions, which when  
4 executed by the one or more processors, causes the one or more processors to perform  
5 the method recited in Claim 17.

1 39. A system comprising:  
2 one or more processors; and  
3 a machine-readable medium carrying one or more sequences of instructions, which when  
4 executed by the one or more processors, causes the one or more processors to perform  
5 the method recited in Claim 18.